

The equivalence between weak and strong purifying selection

Benjamin H. Good¹ and Michael M. Desai¹

¹*Department of Organismic and Evolutionary Biology, Department of Physics, and FAS Center for Systems Biology, Harvard University*

(Dated: October 17, 2012)

Weak purifying selection, acting on many linked mutations, may play a major role in shaping patterns of molecular evolution in natural populations. Yet efforts to infer these effects from DNA sequence data are limited by our incomplete understanding of weak selection on local genomic scales. Here, we demonstrate a natural symmetry between weak and strong selection, in which the effects of many weakly selected mutations on patterns of molecular evolution are equivalent to a smaller number of more strongly selected mutations. By introducing a coarse-grained “effective selection coefficient,” we derive an explicit mapping between weakly selected populations and their strongly selected counterparts, which allows us to make accurate and efficient predictions across the full range of selection strengths. This suggests that an effective selection coefficient and effective mutation rate — not an effective population size — is the most accurate summary of the effects of selection over locally linked regions. Moreover, this correspondence places fundamental limits on our ability to resolve the effects of weak selection from contemporary sequence data alone.

Purifying selection maintains important biological function by purging deleterious mutations and is thought to play a major role in shaping the patterns of molecular evolution in many organisms (Charlesworth, 2012). In principle, these patterns can provide important information about the selective forces operating within a population, and could be used to disentangle this signal from other factors, such as demographic history (Williamson et al., 2005). Yet existing methods are limited by our incomplete understanding of purifying selection on local genomic scales, where many linked sites are potentially selected against.

The action of selection on neighboring sites creates correlations within a genotype that can be difficult to disentangle from each other. Early treatments assumed that these correlations were essentially equivalent to an increase in genetic drift, or a reduction in effective population size, and that the individual sites otherwise evolve independently (Charlesworth, 2009; Hill and Robertson, 1966). Recent studies of these “Hill-Robertson interference” effects have challenged the validity of this assumption (Bustamante et al., 2001; Comeron and Kreitman, 2002; Comeron et al., 2008; Santiago and Caballero, 1998), particularly for the case of weak selection. But without a simple alternative, the effective population size picture continues to dominate much of our qualitative understanding of linked selection and its application to data from natural populations.

Meanwhile, attempts to incorporate linkage more explicitly have been limited to the case where the strength of purifying selection is strong and the number of deleterious polymorphisms is small. In this regime, correlations within genotypes are still highly uncertain, but the distribution of fitnesses within the population can be modeled very precisely. For extremely strong selection,

this leads to the classic background selection picture, in which the apparent size of the population is reduced to the size of the least-loaded class (Charlesworth et al., 1993). More generally, methods based on the structured coalescent framework (Kaplan et al., 1988) lead to improved (though more complicated) analytical predictions (Nicolaisen and Desai, 2012; Walczak et al., 2012), as well as a class of extremely efficient backward-time simulations that can be used to rapidly calculate any quantity of interest (Gordo et al., 2002; Hudson and Kaplan, 1994).

Yet there is increasing evidence that at these local genomic scales, selection is dominated not by a few strongly deleterious polymorphisms, but rather by many more weakly selected mutations that can segregate in the same lineage (Barraclough et al., 2007; Bartolomé and Charlesworth, 2006; Comeron and Kreitman, 2002; Kaiser and Charlesworth, 2008; Loewe and Charlesworth, 2007; Lohmueller et al., 2011; Seger et al., 2010; Subramanian, 2012). In this regime, the strong-selection results break down due to the increased importance of stochastic fluctuations, which can carry some deleterious alleles to intermediate or high frequencies while driving others to extinction. In the extreme case, these fluctuations can sometimes lead to the extinction of the wild-type class and the subsequent fixation of a deleterious allele — an effect known as Muller’s ratchet (Muller, 1964). The complexity of these forces has led to the belief that the dynamics of weak selection are of a fundamentally different character than strong selection, and that a new theoretical picture is required to understand them. Various numerical methods have been devised for this regime, but they either become computationally prohibitive for more than a few selected sites (Barton and Etheridge, 2004; Barton et al., 2004; Krone and Neuhauser, 1997;




FIG. 1 The deterministic prediction for the distribution of fitnesses in the population at mutation-selection balance when $\lambda = U_d/s \approx 3$, and a possible ancestral history for a sample of two individuals.

Neuhauser and Krone, 1997) or require computation time that scales with the size of the population (O’Fallon et al., 2010; Seger et al., 2010), similar to traditional forward-time simulations.

The apparent intractability of the weak selection regime is somewhat paradoxical, given that the limit of *infinitely* weak selection is simply a neutral population. Part of this difficulty arises from the fact that existing coalescent models of purifying selection explicitly track the number of deleterious mutations in each individual. An entirely neutral theory that required similar accounting for neutral mutations would be intractable for many of the same reasons. However, this superficial difficulty obscures a more fundamental aspect of weakly selected mutations: individually, they have a negligible impact on the ancestral process and are indistinguishable from their neutral counterparts, but the accumulation of many such mutations can have a significant effect on the overall diversity of the sample.

In the present work, we exploit this separation of scales to establish a correspondence between the strong and weak selection regimes. By relaxing our definition of neutral and selected mutations and introducing a rescaled *effective selection strength*, we demonstrate an equivalence principle relating the patterns of diversity among populations with differing strengths of selection. For a given population in the weak selection regime, this defines a mapping to a corresponding strong-selection model that captures most of the quantitative features of the original population. The previously developed strong selection results can therefore be extended to provide a single, unified theory valid over the entire range of selective effects, which provides valuable qualitative insights into the net effect of purifying selection.

This correspondence has obvious practical benefits for the analysis of DNA sequence data, since the existing strong-selection techniques can generate efficient predictions across a wide range of parameters, and can potentially form the basis for self-consistent inference of the

underlying selective forces and population sizes. These results have important qualitative implications as well, providing a simple and intuitive alternative to the popular (yet flawed) effective population size picture. Rather, our correspondence suggests that a more natural local quantity is an *effective strength of selection*, defined over some characteristic linkage block. However, the equivalence between strong and weak selection — and the equivalence between weakly selected populations themselves — suggests an inherent limit to our ability to resolve selection pressures from contemporary polymorphism data, especially at the level of individual sites.

I. ANALYSIS

In order to quantify the molecular diversity generated by purifying selection at many linked sites, we confine our attention to a simple and well-studied model in which these effects are known to play a major role. We consider a population of N non-recombining haploid individuals that accumulate neutral mutations at rate U_n and suffer deleterious mutations with a constant multiplicative fitness effect s at rate U_d . We assume that the sequences in the population are well described by an infinite sites model in which each mutation occurs at a unique site in the genome, and we neglect compensatory or otherwise beneficial mutations. In addition, we work in the standard diffusion limit $N \rightarrow \infty$, where the scaled parameters Ns , NU_d , and NU_n are sufficient to determine all quantities of interest.

A. Strong selection

The behavior of this model has been well-characterized when selection against the deleterious alleles is sufficiently strong (we discuss the exact conditions below). In this case, the population reaches a steady state in which the continuous influx of deleterious mutations is balanced on average by the action of selection against them. In the limit $Ns \rightarrow \infty$ where genetic drift can be neglected, the expected fraction of individuals with k deleterious mutations (“fitness class k ”) is given by

$$h_k = \frac{\lambda^k}{k!} e^{-\lambda}, \quad (1)$$

where $\lambda = U_d/s$ parameterizes the relative strength of mutation and selection (Haigh, 1978). An example of this distribution is shown in Fig. 1. As long as Eq. (1) provides a good approximation to the actual stochastic class sizes, the corresponding patterns of diversity are equivalent to a demographically structured *neutral* population, where the h_k are treated as fixed subpopulations and deleterious mutations are recast as migration between the h_k ’s (see Appendix A). This is a special case

of the *structured coalescent* introduced by Kaplan et al. (1988), which traces the ancestry of a sample as it moves through the population fitness distribution (see Fig. 1).

This simplified structured coalescent admits approximate analytical calculations for several simple diversity statistics. These reduce to the standard background selection¹ limit when $Ns \rightarrow \infty$ (Charlesworth et al., 1993), in which the linked neutral diversity is equivalent to an unstructured neutral population with effective population size

$$N_e \approx N e^{-\lambda}. \quad (2)$$

Non-neutral corrections arise for $Ns < \infty$ and require more complicated calculations (Gordo et al., 2002; Nicolaisen and Desai, 2012; Walczak et al., 2012). More importantly in practice, the simplified structured coalescent can be used to efficiently simulate genealogies in a time that scales with the size of the sample rather than the size of the population (Gordo et al., 2002; Hudson and Kaplan, 1994). These coalescent simulations can rapidly generate predictions for many statistics of interest, and could potentially be leveraged to enable full-scale inference of population parameters.

However, this simplified picture is crucially dependent on the assumption that the fitness class sizes (and hence the mutation and coalescence rates) are effectively fixed by the deterministic mutation-selection balance in Eq. (1). In general, genetic drift will cause the actual class sizes to fluctuate around these deterministic predictions, so the validity of our assumption will depend on the severity of these fluctuations. Given our assumption of one-way mutation, there is also a nonzero probability that the least-loaded ($k = 0$) class fluctuates to extinction, allowing one of the deleterious alleles to fix. This effect is known as *Muller’s ratchet*, and it implies that the deterministic mutation-selection balance is stochastically unstable. In the absence of compensatory forces, the entire population will tend to drift toward lower fitness at some small but nonzero rate.

While it is not surprising that these stochastic forces eventually cause significant deviations from Eq. (1), a quantitative characterization of this breakdown is complicated, and the precise answer depends on the quantity of interest. For example, one could examine fluctuations in the fitness class sizes (Neher and Shraiman, 2012), the transition between the so-called “slow” and “fast” regimes of Muller’s ratchet (Gessler, 1995), the breakdown of the background-selection limit (Gordo et al.,

2002), or empirical estimates of divergence from the structured coalescent (see below). Fortunately, most of these definitions of “strong selection” lead to conditions of the form $Ns e^{-\lambda} \gg g(\lambda)$, where $g(\lambda)$ is some slowly growing function of λ . In practice, the simplified condition

$$Ns e^{-\lambda} \gg 1 \quad (3)$$

is generally sufficient to ensure the validity of the structured coalescent for most quantities of interest. We note, however, that even for selection pressures that are traditionally considered to be strong ($Ns \sim 10$), this condition will be violated if the mutation rates are high enough that many of these mutations segregate in the population at the same time ($\lambda \gg 1$). Thus, the strong-selection regime could be more accurately described as a strong-selection/weak-mutation regime.

B. Equivalence Principle for Weak Selection

As selection grows weaker, the deterministic mutation-selection balance provides an increasingly poor estimate of the distribution of fitnesses within the population, as stochastic fluctuations and Muller’s ratchet take on a larger role. It is instructive to consider the extreme limit where $Ns = 0$, when these stochastic forces are strongest. Although we do not normally visualize a neutral population in this way, we could also partition it into “fitness classes” according to the number of mutations in each individual. In this case, the resulting fitness classes fluctuate wildly on coalescent timescales, and Muller’s ratchet “clicks” at rate U_d .

But if Ns is identically zero, this population should also be described by the standard neutral coalescent, which ignores all of these complicated factors. Instead of explicitly tracking the number of mutations in each individual, the neutral coalescent places the entire population within a single fitness class of size N , where fluctuations can be neglected. This simplification arises because in a neutral population, it does not matter which mutations are accounted for by the fitness classes and which accumulate *within* a fitness class, so long as the total mutation rate is preserved. A population with “weakly selected” deleterious mutations of effect $s = 0$ is equivalent to one with “strong selection” ($s' > 0$) if we simultaneously take $U'_d \rightarrow 0$ and $U'_n \rightarrow U_n + U_d$.

This correspondence between weak and strong selection in the neutral limit is admittedly rather trivial, but it suggests that a similar reorganization may hold more generally for weak but non-vanishing Ns . Intuitively, we seek a coarse-grained version of the fitness distribution at some larger scale s' , such that fitness differences less than s' are ignored, but clusters of mutations with cumulative effect s' are treated as a single, large-effect mutation (see Fig. 2). In addition, we wish to choose s'

¹ There is some ambiguity in the literature regarding the term “background selection,” specifically whether it refers to the general effects of purifying selection at linked neutral loci or to the limiting behavior that arises when selection is extremely strong. Here, we use the term in the latter sense as defined by Eq. (2)




FIG. 2 An intuitive picture of the coarse-graining proposed for weakly-selected populations, where the population fitness distribution (only a portion of which is shown here) consists of a large number of fitness classes whose sizes fluctuate considerably. Clusters of several fitness classes are grouped together into larger, *effective* fitness classes separated by fitness \tilde{s} . Deleterious mutations within an effective fitness class are recast as neutral mutations, and only mutations between effective classes are treated as deleterious.

and the reorganized mutation rates U'_d and U'_n in order to mimic (as closely as possible) the patterns of diversity in the original population. It is not clear that this equivalent population should exist *a priori*, and even if it does, the new parameters Ns' , NU'_d , and NU'_n could potentially depend on the underlying parameters Ns , NU_d , and NU_n in some complicated way. Nevertheless, we demonstrate below that an explicit equivalence principle can be obtained from a few simple considerations.

In our neutral example above, we saw that populations with $Ns = 0$ and various combinations of U_d and U_n are equivalent as long as the total mutation rate $NU_{\text{tot}} = NU_d + NU_n$ is preserved. For populations with $Ns > 0$, it is reasonable to expect an additional constraint on the overall scale of selection, which was automatically preserved in the neutral case. This scale is not determined by individual selected mutations, but rather by the emergent distribution of fitnesses within the population. Of course, the fitness distribution is difficult to characterize in weakly-selected populations precisely because of the complicated stochastic effects discussed above. And even if a full solution was available, it is unlikely that the fitness distributions for two different (Ns, NU_d) combinations would exactly coincide. Fortunately, previous studies suggest that the dominant feature of the fitness distribution in the $Ns \rightarrow 0$ limit is the variance σ^2 (Good and Desai, 2012; O'Fallon et al., 2010), which gives a measure of the typical reproductive difference between a random pair of individuals. When the full effects of drift are included, the variance in fitness within the population is given by

$$\sigma^2 = U_d s \left[1 - \frac{R}{U_d} \right], \quad (4)$$

where R is the deleterious substitution rate (Haigh,

1978). We calculate this rate in Appendix B, which completely determines $\sigma^2(Ns, NU_d)$ as a function of the underlying parameters Ns and NU_d .

Thus, we propose an equivalence principle between weakly selected populations, in which the patterns of diversity are equal when

$$\begin{aligned} NU_n + NU_d &= NU'_d + NU'_n, \\ \sigma^2(Ns, NU_d) &= \sigma^2(Ns', NU'_d). \end{aligned} \quad (5)$$

In the simple case where the deterministic approximation $\sigma^2 \approx U_d s$ is valid, we have $U'_d = U_d (s/s')$, which has a natural interpretation in terms of the coarse-grained fitness distribution depicted in Fig. 2. It is important to note that the equivalence defined by Eq. (5) is only *approximately* correct, and it is only valid up to some maximum strength of selection where moments of the fitness distribution other than σ^2 start to become important. This breakdown is unsurprising: when $Nse^{-\lambda} \gg 1$ we must recover the strong selection limit, where it is known that the least-loaded class plays a much larger role than the bulk of the fitness distribution (Charlesworth et al., 1993; Neher and Shraiman, 2012; Nicolaisen and Desai, 2012).

When our equivalence principle holds, Eq. (5) defines an *equivalence class* of populations indexed by a particular value of NU_{tot} and $N\sigma$, in which different underlying parameters Ns , NU_d , and NU_n nevertheless generate the same patterns of diversity. Yet the vast majority of these populations lie well beyond the range of validity of existing methods like the structured coalescent. In order to actually *predict* the patterns of diversity in these populations, we therefore look for a representative population $(N\tilde{s}, N\tilde{U}_d, N\tilde{U}_n)$ within each equivalence class where selection is simultaneously weak enough for our equivalence principle to hold and yet strong enough for these previously developed techniques to be valid.

From our numerical and analytical studies of the structured coalescent (see Appendix C), we have seen that the predictions for the mean pairwise coalescence time $\langle T_2 \rangle / N$ reach a minimum for a particular value of Ns , below which the predictions rapidly diverge from the results of forward-time simulations (see Fig. 3b). It therefore seems reasonable to take $N\tilde{s}$ to be this minimum point, which satisfies

$$\left(\frac{\partial \langle T_2 \rangle}{\partial s} \right)_{\sigma^2} = 0. \quad (6)$$

The subscript denotes that the fitness variance σ^2 is to be held constant when taking the derivative. In Appendix C, we show how this derivative can be calculated using the methods outlined in Walczak et al. (2012). The resulting locus of points yields a *critical line* in the (Ns, NU_d) plane parameterized by the fitness variance $(N\sigma)^2$, as depicted in Fig. 3a. Each point along this critical line corresponds to a coarse-grained model with

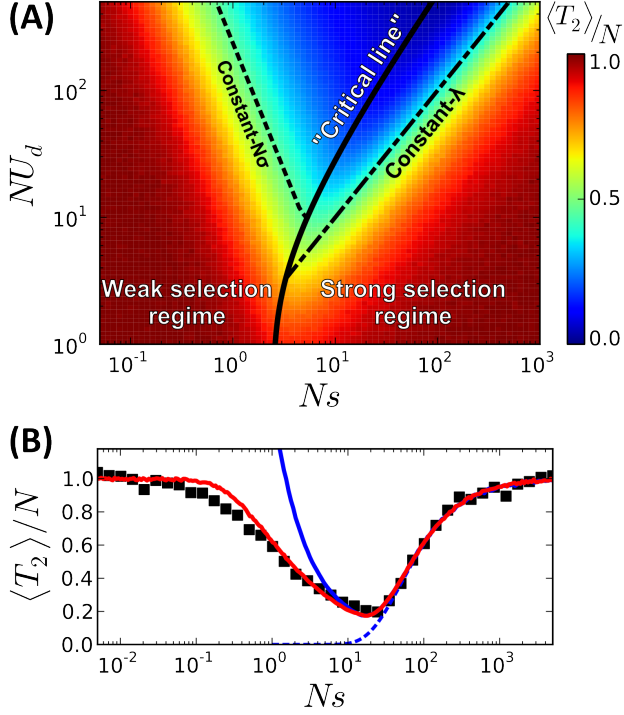


FIG. 3 (A) Predictions for the mean pairwise coalescence time $\langle T_2 \rangle / N$ obtained from structured coalescent simulations of our coarse-grained model. The solid black line denotes the “critical line” ($N\tilde{s}, N\tilde{U}_d$) defined by Eq. (6), while the dashed lines to the left and right denote lines of constant $N\sigma$ and lines of constant λ , respectively. (B) A “slice” of this phase plot at constant $NU_d = 50$. The black squares denote the results of forward-time, Wright-Fisher simulations and our coarse-grained predictions are shown in solid red. For comparison, the solid blue line shows the original structured coalescent predictions, while the dashed blue line shows the background selection approximation $\langle T_2 \rangle \approx Ne^{-\lambda}$.

parameters ($N\tilde{s}, N\tilde{U}_d$) where the structured coalescent is valid, and which can be used to predict the patterns of diversity for all weakly selected populations within that equivalence class. For any particular set of parameters, the corresponding coarse-grained model can be easily calculated from Eqs. (5) and (6), with the help of our expression for σ^2 in Eq. (4).

II. RESULTS

In the previous section, we argued that the patterns of molecular diversity should be equivalent for weakly selected populations with the same total mutation rate and variance in fitness. Within each of these equivalence classes, we have also identified a particular “coarse-grained” population ($N\tilde{s}, N\tilde{U}_d, N\tilde{U}_n$) where previous strong selection methods based on the structured coalescent can be applied. This mapping yields explicit

predictions for various diversity statistics across the full range of selection strengths, an example of which is shown in Fig. 3a for the mean pairwise coalescent time $\langle T_2 \rangle / N$. Parameters that fall to the right of the critical line (depicted by the solid line in Fig. 3a) lie in the strong selection regime where $\langle T_2 \rangle / N$ is directly calculated from the structured coalescent. The vast majority of these points are well-characterized by the background selection limit in Eq. (2), which implies that the level sets of $\langle T_2 \rangle / N$ lie along lines of constant λ (the dashed-dotted line in Fig. 3a). As observed in previous studies, this strong selection equivalence starts to break down near the critical line where the full structured coalescent is required to obtain accurate predictions (Gordo et al., 2002; Walczak et al., 2012). Those parameter values that lie to the left of the critical line are the domain of our coarse-grained theory and the corresponding equivalence along lines of constant $N\sigma$ (depicted by the dashed line in Fig. 3a). We obtain predictions for these populations by applying the structured coalescent to the coarse-grained parameters ($N\tilde{s}, N\tilde{U}_d, N\tilde{U}_n$) calculated using the procedure in Appendix C. Intuitively, this amounts to tracing the line of constant $N\sigma$ back to the corresponding point on the critical line in Fig. 3a.

The accuracy of these predictions rests on two crucial assumptions, which we verify using forward-time, Wright-Fisher simulations (described in Appendix D) for several important and experimentally relevant diversity statistics. First, populations with the same fitness variance $(N\sigma)^2$ should yield similar results for various diversity statistics. Secondly, structured coalescent predictions should agree with the results of forward time simulations along the critical line ($N\tilde{s}, N\tilde{U}_d$).

As we demonstrate in Figs. 4 and 5, both of these assumptions are approximately valid across a large range of parameter values. In Fig. 4a, we plot the mean pairwise coalescent time $\langle T_2 \rangle / N$ obtained from forward-time simulations for a large collection of parameters spanning several orders of magnitude in Ns and NU_d , all of which lie to the left of the critical line in Fig. 3a. The results are organized by their observed fitness variance along the x -axis and colored according to the selection strength of the simulated population. Differently colored points at the same value of $N\sigma$ represent populations with different underlying parameters that fall within the same predicted equivalence class. If our equivalence principle is correct, these colored points should all lie on the same line. In addition, we also plot the structured coalescent predictions for the coarse-grained parameters ($N\tilde{s}, N\tilde{U}_d$) as a function of $N\sigma$, which show good agreement with both forward-time simulations of the critical line (black triangles) and the other populations in each equivalence class.

One of the reasons that the pairwise coalescent time $\langle T_2 \rangle$ plays such a prominent role in earlier studies is that in the standard Hill-Robertson picture, it is equivalent to




FIG. 4 (A) The mean pairwise coalescence time $\langle T_2 \rangle / N$ in the weak-selection regime (top), collated from forward-time, Wright-Fisher simulations at five “slices” of constant $NU_d = 1, 10, 50, 100$ and 300 (similar to Fig. 3b), and the three “slices” of constant $N\sigma = 2, 7, 43$ shown in Fig. 5. Each simulated population is plotted according to its fitness variance $N\sigma$ (averaged over the simulation run) and colored according to the underlying selection strength Ns . In addition, direct forward-time simulations of the critical line ($N\bar{s}, N\bar{U}_d$) are shown as black triangles, while the predictions from the structured coalescent are shown in solid red. (B) A similar plot for the total number of segregating sites S_n in a sample of $n = 100$ individuals, where the total mutation rate is given by $NU_{\text{tot}} = 350$. For comparison, the blue line shows the predictions obtained by assuming independent evolution at different sites, with an effective population size N_e fitted from $\langle T_2 \rangle$ above.

the effective population size that supposedly captures all of the effects of linked selection. Even when $\langle T_2 \rangle$ cannot be predicted analytically, it can be measured from the average heterozygosity at putatively neutral (e.g. synonymous) sites. Both the selected and non-selected sites are then assumed to evolve independently with an effective population size $N_e = \langle T_2 \rangle$. While this intuition works well for predicting the average nonsynonymous heterozygosity [Kaiser and Charlesworth (2008), see Appendix C], several studies have shown that it fails for other statistics that are more sensitive to the correlations produced by linked selection (Comeron and Kreitman, 2002; Comeron et al., 2008; Santiago and Caballero, 1998). In Fig. 4b, we plot the total number of segregating sites S_n in a sample of $n = 100$ individuals in a similar manner as Fig. 4a. We see that even after conditioning on the “correct” reduction in $N_e = \langle T_2 \rangle$, this independent sites assumption significantly underestimates the total diversity

in the sample when $N\sigma > 1$. By contrast, the structured coalescent predictions of the coarse-grained model yield accurate results for the full range of parameters, without needing to fit to the correct $\langle T_2 \rangle$.

In addition to reducing the overall levels of diversity described by statistics like $\langle T_2 \rangle$ and S_n , it is also well-known that purifying selection alters the relative branch lengths in the genealogy of a sample. This distortion is typically measured using the polymorphic site frequency spectrum or one of its derivatives such as Tajima’s D (Tajima, 1989) or Fu and Li’s D (Fu and Li, 1993). When normalized by the total number of singletons, the neutral expectation for the frequencies f_i of the sites polymorphic in i individuals is given by the parameter-free estimate

$$f_i^{\text{neutral}} = i^{-1}. \quad (7)$$

Purifying selection leads to an increase in rare variants ($i \ll n$) compared to this neutral expectation, since deleterious mutations are typically purged before they can drift to appreciable frequencies. Unlike $\langle T_2 \rangle$ or S_n , the site frequency spectrum can be used to detect deviations from neutrality without requiring previous estimates of population size or mutation rate, and is therefore highly useful in the analysis of real sequence data. In Fig. 5, we plot the site frequency spectrum for a sample of $n = 100$ individuals for a range of populations along three particular lines of constant $N\sigma$ and NU_{tot} , but to the left of the critical line in Fig. 3a. Again, we see that populations in the same equivalence class possess very similar frequency spectra, which turn increasingly non-neutral with larger $N\sigma$. We also show the structured coalescent predictions of the coarse-grained model for each value of $N\sigma$, which agree quite well with these forward-time simulations.

However, despite the generally good agreement with forward-time simulations, some small systematic errors remain. For example, the postulated equivalence between populations with the same fitness variance is only approximately true. As can be seen in Figs. 4 and 5, populations that are further from the critical line are generally slightly “less neutral” than their counterparts closer to the critical line, although these differences are often dwarfed by the variation along the critical line itself. In addition, the accuracy of the structured coalescent along the critical line is diminished within a region with low $N\bar{s}$ and $N\bar{U}_d$, which leads us to slightly overestimate overall levels of diversity in this region. These issues are discussed in more detail in Appendix C.

III. DISCUSSION

We have demonstrated an approximate symmetry between the patterns of molecular evolution generated by weak and strong purifying selection. Weakly selected mutations have a negligible impact individually, but a sufficiently large number of these mutations will combine to




FIG. 5 The polymorphic site frequency spectrum for a sample of $n = 100$ individuals, collated from populations with total mutation rate $NU_{\text{tot}} = 350$ and fitness variance $N\sigma \approx 2$ (blue), $N\sigma \approx 7$ (green), and $N\sigma \approx 43$ (red). Symbols denote the results of forward-time simulations (which are shaded from light to dark with decreasing $N\sigma$), while the corresponding structured coalescent predictions are shown as solid lines. For comparison, the prediction for a completely neutral population is depicted by the dashed line. The inset shows the same figure on a log-log scale.

mimic the effects of a single, stronger mutation whose scale is set by the typical fitness differences within the population. This correspondence allows us to import a large body of theory originally developed for the strong selection regime, which we can use to obtain highly accurate predictions across a much broader range of selection strengths than was previously possible.

Our results are consistent with observations from earlier simulation studies of weak selection (Comeron and Kreitman, 2002; Gordo et al., 2002; Kaiser and Charlesworth, 2008; McVean and Charlesworth, 2000; Seger et al., 2010), but our coarse-grained model offers a radically different perspective on the relevant processes that contribute to molecular evolution in this regime. Previous work has argued that virtually all of the deviations from the traditional background selection limit can be attributed to the effects of Muller’s ratchet (Gordo et al., 2002; Seger et al., 2010), as well as to the influence of weak Hill-Robertson interference, where large fluctuations drive weakly selected alleles to intermediate frequencies (Comeron and Kreitman, 2002; McVean and Charlesworth, 2000; Seger et al., 2010). In contrast, our coarse-grained theory includes neither of these complications (aside from the corrections to σ^2) and still captures the quantitative patterns of variation over broad scales. This suggests that variance in ancestral fitness — not fluctuations or the ratchet — is the driving force behind the large-scale patterns of diversity. More complicated stochastic effects may be essential for a first-principles account of linked selection, or for more exotic parameter ranges, but they appear to be of secondary importance

for the quantities and regimes considered here.

Our finding that strong selection methods can be used across the full range of parameters has important practical implications for the inference of selection pressures and population sizes from DNA sequence data. Several previous studies have uncovered evidence for the role of weak purifying selection in natural populations (Barraclough et al., 2007; Betancourt et al., 2009; Loewe and Charlesworth, 2007; Seger et al., 2010), but this evidence is limited by the difficulty in obtaining proper estimates of the population size and selection strength without accounting for the influence of selection at neighboring sites. Previous analyses either ignore linkage altogether (Boyko et al., 2008; Hartl et al., 1994; Keightley and Eyre-Walker, 2007; Tamuri et al., 2012; Williamson et al., 2005) or depend on computationally costly forward-time simulations evaluated over a narrow range of parameters (Kaiser and Charlesworth, 2008; Lohmueller et al., 2011; Seger et al., 2010). Estimates obtained from the first method should be treated cautiously, since the results here and elsewhere demonstrate that the independent-sites assumption is drastically violated when selection at linked sites is common. Although these effects can be treated more rigorously in simulations, our present analysis has uncovered an approximate equivalence or *degeneracy* between populations with weak and strong purifying selection. Equivalent populations are distributed along low-dimensional “ridges” embedded in the larger space of parameters, which can be easily missed when simulating a discrete set of parameter values. The degeneracy identified here arises in a simple model with only three parameters, and it is likely that the number of degeneracies will only increase in more complicated models which involve many more parameters. Thus, our analysis argues for a degree of caution when interpreting the estimates from simulation studies as well, since it can be difficult to determine whether there are other “equivalent” sets of parameters which are equally (or only slightly less) consistent with the data. Ideally, one could combine the knowledge of the degeneracies identified here with the more computationally efficient coalescent simulations to devise a self-consistent inference scheme that utilizes frequency spectrum data, or potentially even phylogenetic reconstruction (Rambaut et al., 2008). A concrete implementation is beyond the scope of the present paper, but this remains an important avenue for future work.

The correspondence between weak and strong selection also suggests an important *qualitative* shift in the interpretation of polymorphism data. It has been known for some time that the effects of linked selection are more complicated than the traditional picture of independent evolution at a reduced effective population size. But without a simple alternative, the concept of a local effective population size — which influences the efficacy of selection and varies along the genome in accordance with the neutral heterozygosity — remains a popular means

of interpreting genomic data (Charlesworth, 2009; Grossmann et al., 2011). In agreement with previous studies, we have provided further evidence that a simple reduction in effective population size does not lead to consistent results for any statistic other than the average heterozygosity π , even after fitting N_e to reproduce the observed reduction in neutral diversity. This leads us to question the ultimate utility of the local effective population size, given that it requires a fit in order to describe only one other property of the data. Rather, our coarse-grained correspondence suggests that a more suitable local quantity is an *effective strength of selection*. In addition to its increased accuracy in capturing the patterns of diversity as demonstrated above, an effective strength of selection is a more natural candidate for a local measure of linked selection, since different parts of the genome are already under varying degrees of selection.

However, we must be careful when interpreting this local effective selection strength due to the degeneracy in the parameter space discussed above. The accuracy of the collapse plots in Figs. 4 and 5 hints at a fundamental resolution limit for inferring the underlying selection pressures from polymorphism data alone, since there is little statistical power to differentiate a weakly selected population from its coarse-grained counterpart (or indeed, any other weakly selected population with the same overall mutation rate and fitness variance). This degeneracy is especially problematic for detecting selection at individual sites, given that the coarse-grained mapping works by reassigning selection from some mutations to others with minimal impact on the overall patterns of diversity.

Of course, our analysis is based on a highly simplified model, and additional work will be required to extend these results to more biologically realistic scenarios. Depending on the particular parameter regime involved, epistasis (Kimura and Maruyama, 1966), finite-site effects (Desai and Plotkin, 2008), and the presence of beneficial or compensatory mutations (Goyal et al., 2012) may all play a larger role than we have assumed here. Particularly questionable is our assumption that all deleterious mutations have the same strength, since it is known that deleterious mutations have a wide distribution of fitness effects in many organisms (Eyre-Walker and Keightley, 2007). Nevertheless, the issues raised here are likely to be a factor in *any* model that includes a sufficiently large number of weakly-selected mutations. At present, there are no analytical descriptions of most of these more complicated scenarios even in the strong-selection regime. Thus while a coarse-grained model of weak selection could be defined in many of these cases, we must first understand the corresponding strong-selection model before coarse-graining can provide useful analytical predictions.

Possibly more problematic for immediate data analysis is our neglect of recombination in the history of the

sample, which limits the direct applications of our theory to asexual organisms, or to mitochondrial DNA or non-recombining regions of the genome in sexual populations. The qualitative issues here remain important in the presence of recombination, since the diversity at each site is influenced by the aggregate selection within some non-recombining neighborhood around it. However, a quantitative extension of these ideas is difficult, since selection and recombination jointly determine the typical linkage scale and the resulting selection regime within that region (Cameron and Kreitman, 2002; Kaiser and Charlesworth, 2008; McVean and Charlesworth, 2000). We note however that recent work by Zeng and Charlesworth (2011) has incorporated finite but nonzero recombination rates into the structured coalescent framework. Thus when an analogous structured coalescent can be defined, our coarse-graining picture will likely be useful for understanding weak selection in these regimes as well. This remains an important avenue for future work.

ACKNOWLEDGMENTS

We thank Richard Neher, Aleksandra Walczak, John Wakeley, Oskar Hallatschek, Lauren Nicolaisen, and Elizabeth Jerison for useful discussions. This work was supported in part by the James S. McDonnell Foundation, the Alfred P. Sloan Foundation, and the Harvard Milton Fund. B.H.G. acknowledges support from a National Science Foundation Graduate Research Fellowship. Simulations in this paper were performed on the Odyssey cluster supported by the Research Computing Group at Harvard University.

REFERENCES

- BARRACLOUGH, T. G., FONTANETO, D., RICCI, C., AND HERNIOU, E. A. 2007. Evidence for inefficient selection against deleterious mutations in cytochrome oxidase i of asexual bdelloid rotifers. *Mol. Biol. Evol.* 24:1952–1962.
- BARTOLOMÉ, C. AND CHARLESWORTH, B. 2006. Evolution of amino-acid sequences and codon usage on the *Drosophila miranda* neo-sex chromosomes. *Genetics* 174:2033–2044.
- BARTON, N. AND ETHERIDGE, A. M. 2004. The effect of selection on genealogies. *Genetics* 166:1115–1131.
- BARTON, N., ETHERIDGE, A. M., AND STURM, A. K. 2004. Coalescence in a random background. *Ann Appl Probab* 14:754–778.
- BETANCOURT, A. J., WELCH, J. J., AND CHARLESWORTH, B. 2009. Reduced effectiveness of selection caused by a lack of recombination. *Current Biology* 19:655–660.
- BOYKO, A. R., WILLIAMSON, S. H., INDAP, A. R., DEGENHARDT, J. D., AND R D HERNANDEZ, E. 2008. Assessing the evolutionary impact of amino acid mutations in the human genome. *PLoS Genet* 4:e1000083.
- BUSTAMANTE, C. D., WAKELEY, J., SAWYER, S., AND HARTL, D. L. 2001. Directional selection and the site-frequency spectrum. *Genetics* 159:1779–1788.

- CHARLESWORTH, B. 2009. Fundamental concepts in genetics: Effective population size and patterns of molecular evolution and variation. *Nat. Rev. Genet.* 10:195–205.
- CHARLESWORTH, B. 2012. The effects of deleterious mutations on evolution at linked sites. *Genetics* 190:5–22.
- CHARLESWORTH, B., MORGAN, M. T., AND CHARLESWORTH, D. 1993. The effect of deleterious mutations on neutral molecular variation. *Genetics* 134:1289–1303.
- COMERON, J. M. AND KREITMAN, M. 2002. Population, evolutionary, and genomic consequences of interference selection. *Genetics* 161:389–410.
- COMERON, J. M., WILLIFORD, A., AND KLIMAN, R. M. 2008. The hill-robertson effect: evolutionary consequences of weak selection and linkage in finite populations. *Heredity* 100:19–31.
- DESAI, M. M. AND PLOTKIN, J. B. 2008. The polymorphism frequency spectrum of finitely many sites under selection. *Genetics* 180:2175–2191.
- ETHERIDGE, A., PFAFFELHUBER, P., AND WAKOLBINGER, A. 2007. How often does the ratchet click? facts, heuristics, and asymptotics, pp. 365–390. In P. M. J. Blath and M. Scheutzow (eds.), *Trends in Stochastic Analysis*. Cambridge University Press.
- EWENS, W. J. 2004. *Mathematical Population Genetics*. Springer-Verlag, New York, second edition.
- EYRE-WALKER, A. AND KEIGHTLEY, P. D. 2007. The distribution of fitness effects of new mutations. *Nat Rev Gen* 8:610–618.
- FU, Y. X. AND LI, W. H. 1993. Statistical tests of neutrality of mutations. *Genetics* 133:693–709.
- GESSLER, D. 1995. The constraints of finite size in asexual populations and the rate of the ratchet. *Genet. Res.* 73:119–131.
- GOOD, B. H. AND DESAI, M. M. 2012. Fluctuations in fitness distributions and the effects of weak linked selection on sequence evolution. arXiv:1210.4044.
- GOOD, B. H., ROUZINE, I. M., BALICK, D. J., HALATSCHKE, O., AND DESAI, M. M. 2012. Distribution of fixed beneficial mutations and the rate of adaptation in asexual populations. *Proc. Natl. Acad. Sci.* 109:4950–4955.
- GORDO, I., NAVARRO, A., AND CHARLESWORTH, B. 2002. Muller’s ratchet and the pattern of variation at a neutral locus. *Genetics* 161:835–848.
- GOYAL, S., BALICK, D. J., JERISON, E. R., NEHER, R. A., SHRAIMAN, B. I., AND DESAI, M. M. 2012. Dynamic mutation-selection balance as an evolutionary attractor.
- GROSSMANN, T. I., WOOLFIT, M., AND EYRE-WALKER, A. 2011. Quantifying the variation in the effective population size within a genome. *Genetics* 189:1389.
- HAIGH, J. 1978. The accumulation of deleterious genes in a population. *Theor Pop Biol* 14:251–267.
- HALLATSCHKE, O. 2011. The noisy edge of traveling waves. *Proc. Natl. Acad. Sci. USA* 108:1783–1787.
- HARTL, D. L., MORIYAMA, E. N., AND SAWYER, S. A. 1994. Selection intensity for codon bias. *Genetics* 138:227–234.
- HIGGS, P. AND WOODCOCK, G. 1995. The accumulation of mutations in asexual populations and the structure of genealogical trees in the presence of selection. *J Math Biol* 33:677–702.
- HILL, W. G. AND ROBERTSON, A. 1966. The effect of linkage on limits to artificial selection. *Genet. Res.* 8:269–294.
- HUDSON, R. R. AND KAPLAN, N. L. 1994. Gene trees with background selection, pp. 140–153. In B. Golding (ed.), *Non-Neutral Evolution: Theories and Molecular Data*. Chapman & Hall, London.
- KAISER, V. B. AND CHARLESWORTH, B. 2008. The effects of deleterious mutations on evolution in non-recombining genomes. *Trends in Genetics* 25:9–12.
- KAPLAN, N., DARDEN, T., AND HUDSON, R. 1988. The coalescent process in models with selection. *Genetics* 198:819–829.
- KEIGHTLEY, P. D. AND EYRE-WALKER, A. 2007. Joint inference of the distribution of fitness effects of deleterious mutations and population demography based on nucleotide polymorphism frequencies. *Genetics* 177:2251–2261.
- KIMURA, M. AND MARUYAMA, T. 1966. The mutational load with epistatic gene interactions in fitness. *Genetics* 54:1303–1312.
- KRONE, S. M. AND NEUHAUSER, C. 1997. Ancestral processes with selection. *Theor Popul Biol* 51:210–237.
- LOEWE, L. AND CHARLESWORTH, B. 2007. Background selection in single genes may explain patterns of codon bias. *Genetics* 175:1381–1393.
- LOHMUELLER, K. E., ALBRECHTSEN, A., LI, Y., KIM, S. Y., KORNELIUSSEN, T., ET AL. 2011. Natural selection affects multiple aspects of genetic variation at putatively neutral sites across the human genome. *PLoS Genetics* 7:e1002326.
- MCVEAN, G. A. T. AND CHARLESWORTH, B. 2000. The effects of hill-robertson interference between weakly selected mutations on patterns of molecular evolution and variation. *Genetics* 155:929–944.
- MULLER, H. J. 1964. The relation of recombination to mutational advance. *Mutat Res* 1:2–9.
- NEHER, R. AND SHRAIMAN, B. I. 2012. Fluctuations of fitness distributions and the rate of muller’s ratchet.
- NEUHAUSER, C. AND KRONE, S. M. 1997. The genealogy of samples in models with selection. *Genetics* 145:519–534.
- NICOLAISEN, L. E. AND DESAI, M. M. 2012. Distortions in genealogies due to purifying selection.
- O’FALLON, B. D., SEGER, J., AND ADLER, F. R. 2010. A continuous-state coalescent and the impact of weak selection on the structure of genealogies. *Mol Biol Evol* 27:1162–1172.
- RAMBAUT, A., PYBUS, O. G., NELSON, M. I., VIBOUD, C., AND *et al* 2008. The genomic and epidemiological dynamics of human influenza a virus. *Nature* 453:617–620.
- ROUZINE, I., BRUNET, E., AND WILKE, C. 2008. The traveling-wave approach to asexual evolution: Muller’s ratchet and the speed of adaptation. *Theor Pop Bio* 73:24–46.
- SANTIAGO, E. AND CABALLERO, A. 1998. Effective size and polymorphism of linked neutral loci in populations under directional selection. *Genetics* 149:2105–2117.
- SEGER, J., SMITH, W. A., PERRY, J. J., HUNN, J., KALISZEWSKA, Z. A., ET AL. 2010. Gene genealogies strongly distorted by weakly interfering mutations in constant environments. *Genetics* 184:529–545.
- SUBRAMANIAN, S. 2012. The abundance of deleterious polymorphisms in humans. *Genetics* 190:1579–1583.
- TAJIMA, F. 1989. Statistical method for testing the neutral mutation hypothesis by dna polymorphism. *Genetics* 123:585–595.
- TAMURI, A. U., DOS REIS, M., AND GOLDSTEIN, R. A. 2012. Estimating the distribution of selection coefficients from phylogenetic data using sitewise mutation-selection models. *Genetics* 190:1101–1115.
- TSIMRING, L., LEVINE, H., AND KESSLER, D. 1996. Rna

virus evolution via a fitness-space model. *Phys. Rev. Lett.* 90:088103.

WALCZAK, A. M., NICOLAISEN, L. E., PLOTKIN, J. B., AND DESAI, M. M. 2012. The structure of genealogies in the presence of purifying selection: a fitness-class coalescent. *Genetics* 190:753–779.

WILLIAMSON, S. H., HERNANDEZ, R., FLEDEL-ALON, A., ZHU, L., NIELSEN, R., AND BUSTAMANTE, C. D. 2005. Simultaneous inference of selection and population growth from patterns of variation in the human genome. *Proc Nat Acad Sci* 102:7882–7887.

ZENG, K. AND CHARLESWORTH, B. 2011. The joint effects of background selection and genetic recombination on local gene genealogies. *Genetics* 189:251–266.

Appendix A: The Structured Coalescent

Our theoretical predictions for the diversity statistics in the main text are calculated within the structured coalescent framework, which provides an explicit probabilistic model of the genealogy of a sample from the population. In its most general form, the structured coalescent extends the neutral Kingman coalescent to incorporate arbitrary time-dependent (and possibly stochastic) demographic structure, but this has proven difficult to implement in practice. In the present work, we therefore focus our attention on a particularly simple special case, where the relevant demographic structure is the division of the population into constant fitness classes attained at mutation-selection balance (Hudson and Kaplan, 1994). As such, this simplified structured coalescent is fundamentally a strong selection result.

In the analysis that follows, we work in the standard coalescent limit $N \rightarrow \infty$, where the scaled parameters Ns , NU_d , and NU_n are sufficient to characterize the population, and we measure time in units of N generations. As mentioned above, we assume that the distribution of fitnesses within the population is given by the deterministic mutation-selection balance in Eq. (1) in the main text, and we neglect fluctuations in the class sizes.

We wish to characterize the possible genealogical histories of a sample of n individuals drawn from the population. In a random sample, these individuals come from fitness classes k_1, \dots, k_n drawn from the population fitness distribution, which implies that

$$k_1, \dots, k_n \stackrel{i.i.d.}{\sim} \text{Poisson}(\lambda = NU_d/Ns). \quad (\text{A1})$$

We then trace the genealogy back to the most recent common ancestor of the sample. At any given instant, three types of ancestral events can occur:

1. An individual can experience a neutral mutation at rate

$$N \cdot \left(\frac{NU_n}{N} \right) = NU_n. \quad (\text{A2})$$

2. An individual in class $k_i > 0$ can experience a deleterious mutation (thus transferring it to class $k_i - 1$) at rate

$$N \cdot \left(\frac{U_d N h_{k_i-1}}{N h_{k_i}} \right) = N s k_i \quad (\text{A3})$$

3. Two individuals in the same fitness class $k = k_i = k_j$ can coalesce to a single individual at rate

$$N \cdot (N h_k) \cdot \left(\frac{1}{N h_k} \right)^2 = \frac{1}{h_k} \quad (\text{A4})$$

These events are competing Poisson processes, which implies that the time to the next event is exponentially distributed with mean equal to the reciprocal of the sum of the rates of all possible events. The event itself is then drawn randomly from the pool of possible events, each weighted by its corresponding rate. This process continues until the sample has coalesced into a single lineage, which is the most recent common ancestor of the sample.

Thus, for a given set of parameters Ns , NU_d , and NU_n , the distribution of genealogies and mutation events for a sample of n individuals is completely specified. The distribution of any particular diversity statistic can be straightforwardly obtained by averaging over the distribution of genealogies. In practice, evaluating these averages analytically can be difficult for all but the simplest statistics (Walczak et al., 2012). Instead, it is often easier to use the procedure outlined above to implement backward-in-time simulations that sample genealogies from this ancestral process (Gordo et al., 2002). These coalescent simulations are extremely computationally efficient compared to their ordinary forward-time counterparts, since we only need to simulate ancestral events for the sample as opposed to simulating the entire population at each generation. A copy of our implementation in C is available upon request.

Appendix B: Fitness variance under weak selection

In the main text, we postulate an equivalence principle between weakly selected populations with the same variance in fitness, which we calculate here. When selection is strong, σ^2 can easily be calculated from the mutation-selection balance in Eq. (1) in the main text, and we find that

$$\sigma^2 = U_d s. \quad (\text{B1})$$

However, we wish to apply these results precisely in the region where the deterministic mutation-selection balance becomes unreliable, so we would like to generalize this calculation to include the full effects of drift as

$Ns \rightarrow 0$. A standard calculation shows that more generally, we have

$$\sigma^2 = U_d s \left(1 - \frac{R}{U_d}\right), \quad (\text{B2})$$

where R is the substitution rate of deleterious mutations (Etheridge et al., 2007; Good and Desai, 2012; Haigh, 1978; Higgs and Woodcock, 1995). This reduces to the ordinary strong-selection result in Eq. (B1) when $Nse^{-\lambda} \gg 1$, but we now focus on the regime where R is not necessarily small compared to U_d . Fortunately, when R is on the order of U_d , this substitution rate is equivalent to the rate of Muller's ratchet in the so-called "fast-ratchet" regime (Gessler, 1995), where many of the complicated stochastic aspects of the ratchet analyzed in Neher and Shraiman (2012) can be neglected and results from traveling wave theory (Good et al., 2012; Hallatschek, 2011; Rouzine et al., 2008) can be applied.

We calculate the rate of deleterious substitutions using the tunable constraint framework introduced in Hallatschek (2011), which modifies the standard Wright-Fisher stochastic dynamics in order to make R easier to calculate. These predictions for the substitution rate have been shown to agree with ordinary Wright-Fisher simulations in several regimes of positive selection (Good et al., 2012; Hallatschek, 2011), but they have yet to be directly applied to the purifying selection regime studied here.

We introduce two new quantities $f(x)$ and $w(x)$, which respectively correspond to the population density and the fixation probability of new mutants at relative fitness x . In the tunable constraint framework, these are related to the population size and substitution rate through the system of equations

$$sR\partial_x f(x) = xf(x) - f(x)w(x) + U_d[f(x+s) - f(x)], \quad (\text{B3})$$

$$-sR\partial_x w(x) = xw(x) - w(x)^2 + U_d[w(x-s) - w(x)], \quad (\text{B4})$$

and the normalization conditions

$$1 = \int_{-\infty}^{\infty} f(x) dx, \quad \frac{\nu}{N} = \int_{-\infty}^{\infty} f(x)w(x) dx, \quad (\text{B5})$$

where ν is the variance in offspring number (equal to unity in the Wright-Fisher model). This system of equations can in principle be solved to obtain R as a function of N , s , and U_d , but in their current form, these equations are difficult to solve (even numerically) due to the delay terms $f(x+s)$ and $w(x-s)$ in the differential equations. Thus, we turn to an approximate solution.

Since we are focused on a regime where selection is weak, it seems reasonable to try a Taylor expansion of these delay terms in powers of s :

$$f(x+s) = f(x) + s\partial_x f(x) + \frac{1}{2}\partial_x^2 f(x) + \dots \quad (\text{B6})$$

In order to obtain a non-trivial solution for $w(x)$, we must include at least the second order term in this expansion, after which Eqs. (B3) and (B4) can be rewritten in the form

$$0 = xf(x) - w(x)f(x) + \left(\frac{U_d s^2}{2}\right) \partial_x^2 f(x) + U_d s \left[1 - \frac{R}{U_d}\right] \partial_x f(x), \quad (\text{B7})$$

$$0 = xw(x) - w(x)^2 + \left(\frac{U_d s^2}{2}\right) \partial_x^2 w(x) - U_d s \left[1 - \frac{R}{U_d}\right] \partial_x w(x). \quad (\text{B8})$$

Note that this is essentially the same model analyzed in (Hallatschek, 2011), with the parameters D and v of that work corresponding to $U_d s^2/2$ and $U_d s(1 - R/U_d)$ here. A similar analysis can therefore be applied in the present case. We first rescale the relative fitness x by introducing the new coordinate

$$\chi = x \left(\frac{U_d s^2}{2}\right)^{-1/3}, \quad (\text{B9})$$

Since $w(x)$ and $f(x)$ currently have the units of fitness and inverse fitness, respectively, we must rescale these as well:

$$\begin{aligned} \tilde{f}(x) &= \left(\frac{U_d s^2}{2}\right)^{1/3} f(x), \\ \tilde{w}(x) &= \left(\frac{U_d s^2}{2}\right)^{-1/3} w(x). \end{aligned} \quad (\text{B10})$$

In terms of these rescaled variables, our system of equations can be written in the compact form

$$0 = \partial_\chi^2 \tilde{f} + \alpha \partial_\chi \tilde{f} + \chi \tilde{f} - \tilde{f} \tilde{w}, \quad (\text{B11})$$

$$0 = \partial_\chi^2 \tilde{w} - \alpha \partial_\chi \tilde{w} + \chi \tilde{w} - \tilde{w}^2, \quad (\text{B12})$$

$$1 = \int_{-\infty}^{\infty} \tilde{f}(\chi) d\chi, \quad (\text{B13})$$

$$\frac{1}{\beta} = \int_{-\infty}^{\infty} \tilde{f}(\chi) \tilde{w}(\chi) d\chi, \quad (\text{B14})$$

where we have introduced the two parameters

$$\alpha = (4\lambda)^{1/3} \left(1 - \frac{R}{U_d}\right), \quad \beta = \frac{Ns}{\nu} \left(\frac{\lambda}{2}\right)^{1/3}. \quad (\text{B15})$$

From inspection, we can immediately see that

$$\tilde{f}(\chi) \propto e^{-\alpha\chi} \tilde{w}(\chi) \quad (\text{B16})$$

is a solution to Eq. (B11), which allows us to eliminate \tilde{f} entirely and yields the simplified system

$$\begin{aligned} 0 &= \partial_\chi^2 \tilde{w} - \alpha \partial_\chi \tilde{w} + \chi \tilde{w} - \tilde{w}^2, \\ \beta &= \frac{\int_{-\infty}^{\infty} e^{-\alpha\chi} \tilde{w}(\chi) d\chi}{\int_{-\infty}^{\infty} e^{-\alpha\chi} \tilde{w}(\chi)^2 d\chi}, \end{aligned} \quad (\text{B17})$$




FIG. 6 The variance in fitness within the population as a function of the selection strength, for $NU_d = 5$ (blue), $NU_d = 50$ (green), and $NU_d = 500$ (red). Symbols denote the results of forward-time simulations with $N = 10^5$ averaged over 300 independent runs, while the solid lines are the predictions obtained from Eqs. (B2) and (B23). The dashed lines give the asymptotic behavior for $Ns \rightarrow 0$ and $Ns \rightarrow \infty$.

where $\tilde{w}(\chi)$ is subject to the boundary conditions $\tilde{w}(\chi) \rightarrow 0$ as $\chi \rightarrow -\infty$ and $\tilde{w}(\chi) \rightarrow \chi$ as $\chi \rightarrow \infty$. Thus, numerical solution of the boundary value problem (for instance, using Matlab's `bvp4c` function) and subsequent numerical integration of the resulting \tilde{w} allows us to calculate β as a function of α ,

$$\beta = g(\alpha), \quad (\text{B18})$$

where g is independent of any of the evolutionary parameters. A subsequent inversion of this relation yields an expression for $1 - R/U_d$ as a function of Ns and λ :

$$1 - \frac{R}{U_d} = \left(\frac{1}{4\lambda}\right)^{1/3} g^{-1} \left[\frac{Ns}{\nu} \left(\frac{\lambda}{2}\right)^{1/3} \right]. \quad (\text{B19})$$

For general values of α , this function $g(\alpha)$ must be calculated numerically using the approach outlined above (an implementation in Matlab is available upon request). However, we can obtain simple analytical formulae in two limiting cases. The limit of large α has been studied extensively in previous work (Goyal et al., 2012; Halatschek, 2011; Tsimring et al., 1996), and we find that $\log g(\alpha) \sim \alpha^3$. In the limit $\alpha \rightarrow 0$, the differential equation for \tilde{w} reduces to the parameter-free form

$$0 = \partial_\chi^2 \tilde{w} + \chi \tilde{w} - \tilde{w}^2, \quad (\text{B20})$$

whose solution can be reasonably well-approximated in this limit by

$$\tilde{w}(\chi) \approx \begin{cases} \chi & \text{if } \chi > 0, \\ 0 & \text{else.} \end{cases} \quad (\text{B21})$$

The normalization integrals in Eq. (B17) can then therefore be approximated by Γ -functions, and we find that

$g(\alpha) \approx \frac{\alpha}{2}$. In terms of the underlying evolutionary parameters, this implies that

$$\sigma^2 \approx \left(\frac{NU_d}{\nu}\right) s^2 \quad (\text{B22})$$

as $Ns \rightarrow 0$, which agrees with an analogous calculation using the neutral coalescent.

As a slight technical aside, we note that the left-hand side of Eq. (B19) by definition cannot be greater than one, since the deleterious mutations cannot accumulate at a negative rate. Nevertheless, as Ns increases the right-hand side of Eq. (B19) eventually becomes larger than one (particularly in the large α limit discussed above). This likely indicates a breakdown of the mutational diffusion approximation we assumed in Eq. (B6), or possibly a breakdown in the applicability of the tunable constraint framework in general. In order to maintain sensible results, we therefore take

$$1 - \frac{R}{U_d} = \min \left\{ \left(\frac{1}{4\lambda}\right)^{1/3} g^{-1} \left[Ns \left(\frac{\lambda}{2}\right)^{1/3} \right], 1 \right\}. \quad (\text{B23})$$

In Fig. 6, we compare these predictions with the results of forward-time simulations for a representative sample of parameters. The agreement is generally quite good (certainly better than the naive asymptotics alone), although there are some small systematic disagreements. In particular, we tend to slightly overestimate the variance in fitness at the point where it starts to deviate from the deterministic asymptote $\sigma^2 = U_d s$, and we tend to slightly underestimate it during the transition to the neutral asymptote $\sigma^2 = NU_d s^2$.

Appendix C: The Coarse-Grained Model

In the main text, we proposed a theory of weak purifying selection, which was based on two underlying assumptions:

1. Weakly selected populations with the same NU_{tot} and $N\sigma$ form an equivalence class in terms of the patterns of diversity they contain.
2. Within this equivalence class, there exists a population with parameters $(N\tilde{s}, N\tilde{U}_d, N\tilde{U}_n)$ where the strength of selection is strong enough that the structured coalescent is valid.

This allows us to generate predictions for any weakly selected population by identifying the corresponding “coarse-grained” population and applying the structured coalescent.

1. Finding the coarse-grained parameters

For a particular population with parameters Ns , NU_d , and NU_n , we first identify the corresponding equiva-

lence class by calculating $N\sigma$ using the formulae in Appendix B. Then, as described in the main text, the corresponding coarse-grained population can be found by minimizing $\langle T_2 \rangle / N$ as a function of Ns , with $N\sigma$ held constant. This coarse-grained population is by definition in the strong selection regime, so $\langle T_2 \rangle / N$ can be calculated analytically using the results in (Walczak et al., 2012), which state that

$$\frac{\langle T_2 \rangle}{N} = \sum_{k_1=0}^{\infty} \sum_{k_2=k_1}^{\infty} \sum_{k_c=0}^{k_1} t(k_1, k_2, k_c) p(k_1, k_2, k_c). \quad (C1)$$

The function $t(k_1, k_2, k_c)$ is the mean coalescent time for two individuals in classes k_1 and k_2 that coalesce in class k_c , and is given by

$$t(k_1, k_2, k_c) = \frac{h_{k_c}}{1 + 2Ns k_c h_{k_c}} \quad (C2)$$

in the special case that $k_1 = k_2 = k_c$ and

$$t(k_1, k_2, k_c) = \sum_{j=1}^{k_1+k_2-2k_c} \frac{j(-1)^{j-1}}{(j+2k_c)^2} \binom{k_1+k_2-2k_c}{j} \times \binom{k_1+k_2}{2k_c} \left[\frac{1}{Ns} + \frac{(j+2k_c)h_{k_c}}{1+2Ns k_c h_{k_c}} \right] \quad (C3)$$

otherwise. The function $p(k_1, k_2, k_c)$ is the probability of sampling two individuals from classes k_1 and k_2 that coalesce in class k_c , and is given by

$$p(k_1, k_2, k_c) = \frac{\binom{k_1}{k_c} \binom{k_2}{k_c} \binom{k_1+k_2}{2k_c}^{-1} h(k_1, k_2)}{1 + 2Ns k_c h_{k_c}} \times \prod_{j=k_1}^{k_c-1} \left[1 - \frac{\binom{k_1}{j} \binom{k_2}{j} \binom{k_1+k_2}{2j}^{-1}}{1 + 2Ns j h_j} \right] \quad (C4)$$

where

$$h(k_1, k_2) = \begin{cases} h_{k_1}^2 & \text{if } k_1 = k_2, \\ 2h_{k_1} h_{k_2} & \text{else.} \end{cases} \quad (C5)$$

In addition, when Eq. (C1) is valid, the fitness variance is related to Ns and NU_d through the simple relation

$$(N\sigma)^2 \approx (NU_d)(Ns). \quad (C6)$$

Putting these two facts together, we see that $N\tilde{s}$ is given by the root of the equation

$$0 = \frac{\partial \langle T_2 \rangle}{\partial (Ns)} - \frac{2\lambda}{Ns} \frac{\partial \langle T_2 \rangle}{\partial \lambda}, \quad (C7)$$

where we set $\lambda = (N\sigma/Ns)^2$ after taking the derivative. Once we have determined $N\tilde{s}$, the corresponding mutation rates $N\tilde{U}_d$ and $N\tilde{U}_n$ are given by

$$N\tilde{U}_d = \frac{(N\sigma)^2}{N\tilde{s}}, \quad (C8)$$

$$N\tilde{U}_n = NU_n + (NU_d - N\tilde{U}_d).$$




FIG. 7 The mean pairwise coalescence time as a function of the selection strength, for fixed $N\sigma \approx 2$ (blue), $N\sigma \approx 7$ (green), and $N\sigma \approx 43$ (red). Symbols denote the results of forward-time simulations, while the solid lines show the predictions from the original structured coalescent. The dotted lines denote the corresponding points $(N\tilde{s}, N\tilde{U}_d)$ on the critical line, which are utilized by our coarse-grained theory.

Thus, we have constructed an explicit mapping between underlying parameters (Ns, NU_d, NU_n) and those of the corresponding coarse-grained model $(N\tilde{s}, N\tilde{U}_d, N\tilde{U}_n)$ where we can apply the structured coalescent. A copy of our implementation in Python is available upon request.

2. Deviations from the coarse-grained predictions

This procedure enables us to generate predictions for any set of parameters Ns , NU_d , and NU_n , but the validity of these predictions depends on the validity of the underlying assumptions (1) and (2) stated above. Figure 3 (main text) shows that these assumptions are *approximately* true over a broad parameter regime, but some small systematic deviations are observed. In order to examine these deviations in more detail, we focus on a narrower (yet still representative) set of parameters.

In the same manner as Figure 4 in the main text, we simulate three lines of constant $N\sigma$ and NU_{tot} [calculated from Eqs. (B2) and (B23)] for $N\sigma \approx 2$, $N\sigma \approx 7$, and $N\sigma \approx 43$. Results for the mean pairwise coalescent time are shown in Fig. 7. In all three cases, we observe a characteristic “hockey-stick” shape, in which a region of rapidly varying $\langle T_2 \rangle / N$ sharply transitions to a region of significantly reduced variation. This abrupt transition coincides with the minimum of the structured coalescent predictions for $\langle T_2 \rangle / N$, and hence with the boundary of the strong selection regime.

If assumption (1) is exactly satisfied, the points to the left of this boundary should all have the same value. We see that this is true to a good degree of approximation, in the sense that the remaining variation in these points




FIG. 8 The mean pairwise heterozygosity (top) and the total number of segregating sites for a sample of $n = 100$ individuals (bottom) as a function of selection strength for the lines of constant $N\sigma$ shown in Fig. 7, with $NU_{\text{tot}} = 350$. Symbols denote the results of forward-time simulations, while the solid lines show the predictions from the effective population size assumption, fitted to the simulated values of $\langle T_2 \rangle$ in Fig. 7. The dotted lines denote the corresponding points $(N\tilde{s}, N\tilde{U}_d)$ on the critical line, which are utilized by our coarse-grained theory.

is much less than the variation *between* the lines with different $N\sigma$. Nevertheless, there remains a slight downward trend along these lines of constant $N\sigma$ as selection grows weaker, which corresponds to these points being slightly “less neutral” than we would predict using our equivalence class. In Fig. 8 we plot the mean pairwise heterozygosity and the total number of segregating sites for these lines as well. We observe a similar “hockey stick” shape, although the deviation from assumption (1) is slightly stronger than we observed for $\langle T_2 \rangle / N$.

Our coarse-grained theory also depends on the validity of assumption (2), which requires that the structured coalescent predictions along the critical line should match forward-time simulations without any further modifications. Again, while Figs. 4 and 5 in the main text show that this is generally true, we do observe some systematic deviations, especially for points where both $N\tilde{s}$ and $N\tilde{U}_d$ are small. We can examine this regime more closely in Fig. 9, which plots the pairwise coalescent time as a function of the selection strength for constant $NU_d = 1$.

We observe that at these low values of $\lambda = NU_d / Ns$,




FIG. 9 The mean pairwise coalescent time (top) and the total tree length for a sample of $n = 100$ individuals (bottom) as a function of selection strength, for fixed $NU_d = 1$. Symbols denote the results of forward-time simulations, and the black dashed lines give approximate 95% confidence intervals. The solid blue line shows the predictions from the original structured coalescent, while the dashed blue line is the corresponding background selection approximation.

the structured coalescent overestimates the characteristic minimum value of $\langle T_2 \rangle$, although it gets the location more or less correct. Typically, these low values of λ are associated with extremely strong selection pressures, so that the classic background selection approximation is valid. However, we see that near this minimum – which represents the maximum deviation from neutrality for this level of mutation – neither background selection nor the structured coalescent gives the correct result. Our coarse-grained theory can therefore do no better.

Further study of this low NU_d and low Ns region may shed light on the interactions between stochastic fluctuations and the structured coalescent framework, and could offer insight on how to incorporate first order corrections for these effects. However, these deviations are generally small, and for such low values of Ns and NU_d the population is nearly neutral anyway. In addition, small discrepancies in this relatively narrow region of parameter space are unlikely to matter much for practical purposes, since the patterns of diversity observed in actual populations are likely to be dominated by larger deviations from neutrality attained at larger values of Ns and NU_d .




FIG. 10 The distribution of neutral pairwise heterozygosity for a population with $Ns = 5$ (left) and $Ns = 0.25$ (right), with $NU_d = NU_n = 50$. Black lines denote the results of forward-time simulations and the red lines show the structured coalescent predictions for our coarse-grained theory. For comparison, the blue lines show the predictions from the standard effective population size picture, with N_e fitted from the mean of the forward-time distribution.




FIG. 11 The relative variance in pairwise coalescent time as a function of $N\sigma$ for the same populations as Figure 3 in the main text. Again, symbols denote the results of forward-time simulations (black dotted lines show approximate 95% confidence intervals), while the solid red line gives the structured coalescent predictions of our coarse-grained theory. For comparison, the predictions from the standard effective population size picture are shown as a solid blue line.

3. Comparison with the effective population size picture

The diversity statistics in Fig. 8 allow us to make a more detailed comparison with the effective population size picture that is typically used to interpret such data. In this case, N_e can be measured exactly by fitting to the corresponding $\langle T_2 \rangle$ values in Fig. 7, and the predictions for π and S_n (shown as solid lines in Fig. 8) follow from assuming that the individual sites otherwise evolve independently at this reduced effective population size. We see that while this assumption leads to excellent agreement for π (as observed previously in (Kaiser and Charlesworth, 2008)), it drastically underestimates S_n , which depends more sensitively on the genealogical distortions caused by purifying selection.

We can find even larger discrepancies by focusing on the *distributions* of certain statistics, — particularly those related to the neutral diversity — which take on an extremely simple form in the effective population size picture. For example, under this assumption the distribution of neutral pairwise heterozygosity is predicted to follow a geometric distribution with mean $2N_e U_n$. The

most likely value is $\pi_n = 0$, and the probability of larger values decreases monotonically with increasing π_n . In Fig. 10, we plot the distribution of neutral heterozygosity for two populations in the weak selection regime. We see that this effective population size assumption fails to capture the qualitative features of the distribution, despite being fitted to the correct mean value. This effect is exaggerated closer to the critical line, where the distribution develops a strong peak at a nonzero value of π_n resulting from a corresponding peak in the distribution of T_2 .

We can quantify this peaked nature of the distribution over a broader range of parameters by looking at the variance in the pairwise coalescent time. Under the effective population size assumption, T_2 follows an exponential distribution with mean N_e . This implies that the ratio of the variance and the mean is given by

$$\frac{\text{Var}(T_2)}{\langle T_2 \rangle^2} = 1, \quad (\text{C9})$$

independent of N_e or any of the other parameters. In Fig. 11, we measure this statistic for each of the populations in Fig. 4 (main text) and construct an analogous collapse plot. Again, our equivalence principle is highly accurate, and our coarse-grained predictions from the structured coalescent quantitatively describe these effects of linked selection that are not even qualitatively captured by the effective population size picture.

Appendix D: Forward-time Simulations

We validate several of our key approximations in the main text by comparing our theoretical predictions with the results of forward-time, discrete-generation simulations similar to the standard Wright-Fisher model (Ewens, 2004). These simulations begin with a clonal population of N individuals, and in each subsequent generation the population undergoes a selection step followed by a mutation step. In the selection step, each lineage (i.e. unique genotype) is assigned a new size from a Poisson distribution with mean

$$\lambda_i = C(1 + x_i)n_i, \quad (\text{D1})$$

where n_i is the current size of the lineage, x_i is its fitness relative to the population average, and $C = N/\sum_i n_i$ is a normalization constant chosen to ensure that the total population size remains close to N . In the mutation step, each individual mutates with probability $U_d + U_n$, and if it does, the new mutation is deleterious with probability $U_d/(U_d + U_n)$. This process is continued for a sufficiently long period of time that the population reaches the steady-state mutation-selection balance introduced in the main text, and several population-wide coalescence events have occurred. A copy of our implementation in C is available upon request.